EXHIBIT S

Arista Response to Sourcing Event for LEAF
~
Addendum A: Arista Executive Summary

PREPARED FOR:

at&t

November 2, 15

PROVIDED BY:

ARISTA

Arista Networks, Inc.

Wayne Stickle
Director of Sales, AT&T
Austin, TX

Office: (512) 263-0571
Mobile: (512) 656-6373
Email: Wayne@Arista.com

Table of Contents

List of Figures

List of Tables

List of Videos

1 Executive Summary

Deploying modern data center equipment in a service provider environment requires a new level of agility. Next generation IP services require an ultra-scalable network infrastructure, which quickly adapts new technology while enabling higher density 10/25/40/50/100GbE switched Ethernet services.

Arista's response to this initiative will highlight and demonstrate distinguishable ROI in these key areas;

- **Reduced Ethernet tax:** ability to deploy fewer devices with greater port density, thus reducing power, space and data center overhead.
- **Standard based solution:** The proposal will based on open standards, ensuring AT&T are not locked into a technology cul-de-sac, allowing them to innovate continuously throughout the life-span of the deployment.
- **Increased Business Agility**: improved time to market for new services by utilizing zero touch provisioning (ZTP), zero touch recovery (ZTR) to enable faster implementation through integration with orchestration and automation tools.
- **Reduced Operational Cost**: ability to manage thousands of devices through seamless integration with existing network management systems, with the ability to deploy an open programmable + API enablement infrastructure. Arista EOS single binary image across all platforms enables more efficient code certification and release management.

1.1 Commitment to Merchant Silicon

Of the leading networking manufacturers, Arista is only one fully committed to solely using commercial full-custom silicon from multiple suppliers across its portfolio. This market is highly competitive and diversified and it ensures Arista can choose from a range of choices for cutting edge, best-of-breed silicon, and avoid having to rely on single-source approaches.

Prior to Arista Networks, no other networking company has driven its product portfolio towards a path that follows Moore's Law in a way only previously experienced in general computing. With single-chip designs now providing the equivalent of 128 10GbE ports in a 1RU form factor and plans in 2015 to double this density to 256 10GbE ports in a single-chip design, Arista is finally changing the game and the economics of Ethernet in the data center. Our customers now have a viable partner to rely on for unmatched density, favorable environmentals, high performance, and cost.

Arista only builds best of breed data center switches; from low-latency, high density top-of-rack, to the highest performing carrier grade, modular chassis switch in the market. This means that AT&T can trust Arista to deliver market leading platforms at Moore's Law cadence, with the highest port density, lowest form factor and most efficient power footprint in the market today; ensuring AT&T always remains at the cutting edge of technology and innovation.

1.2 Open Standards

Arista's commitment to open standards based, highly reliable, data center networking will provide AT&T with the flexibility and resiliency required, while simplifying the overall infrastructure and minimizing operational costs. Our ability to work within the data center eco-system and partner with market leading vendors in SDN, storage,

security, virtualization, management and automation, ensures AT&T can continue to adopt best of breed solutions without fear of being "locked-in" or being driven down a path which may not be best suited to the ever changing needs of the network and IT. With a proven track record of interoperability with the many disparate vendors provided components, Arista can interoperate or replace any and all components of the switching infrastructure.

In summary, we believe that by partnering with Arista, AT&T will realize significant cost benefits in operational efficiencies, CapEx expenditure, power, and hardware footprint.

1.3 LEAF Switch Hardware

The Leaf layer within the topology would act as the first-hop gateway for any directly connected hosts within a rack. The hosts residing within one or multiple IP subnets, and connecting to the Leaf via a trunk (802.1Q) or an access link. The IP subnets would be announced via the IGP to the Spine to allow visibility form the neighboring Leaf nodes

1.3.1 MLAG topology

Within the rack there will be a requirement to dual home hosts and appliances for resilience purposes to the Leaf layer, while maintaining active-active Layer 2/3 forwarding for optimal bandwidth utilization and performance.

To achieve this requirement all Arista switches support Multi-Chassis Link Aggregation (MLAG). The MLAG technology allows two physical Arista switches to appear as a single logical switch called an MLAG Domain. Hosts and third-party switches connect to the MLAG Domain using a standard LACP (static, passive or active) port-channel with the links of the port-channel split across the two physical switches to provide resilient active/active Layer 2 forwarding. By using the LACP protocol the technology is transparent to the attached hosts, allowing the flexibility to connect Arista switches, hosts and third-party switches.

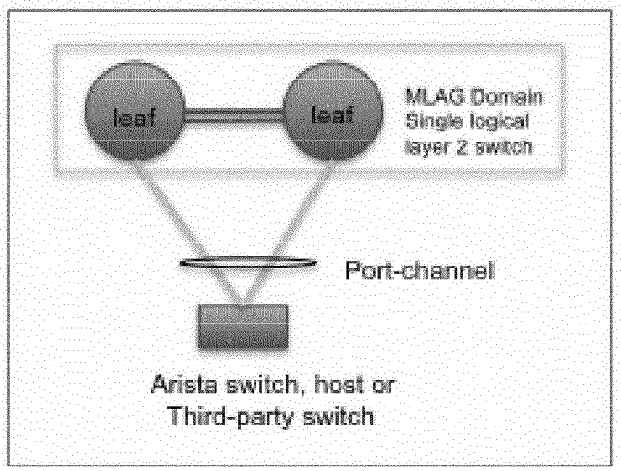

Figure 1. Diagram of an Arista MLAG topology

An MLAG Domain is formed, by interconnecting a pair of Leaf switches via a dedicated port-channel called the MLAG peer link. This peer link is then used to synchronize MAC address tables and port-channel status between the two physical switches

Traffic forwarding within the MLAG Domain always follows the optimal path. Northbound traffic from the host, which would be load-balanced based on the host's own algorithm, would be forwarded by the first switch to

received the traffic. Traffic Southbound, received from the Spine, would be forwarded down the local link of the switch(es) directly connected to the end host. Thus under-normal conditions, no data traffic traverses the peer-link, it is only under a failure condition that the peer link is used to pass data traffic. For this reasons the peer link should be provisioned predominately for resiliency reasons rather than bandwidth.

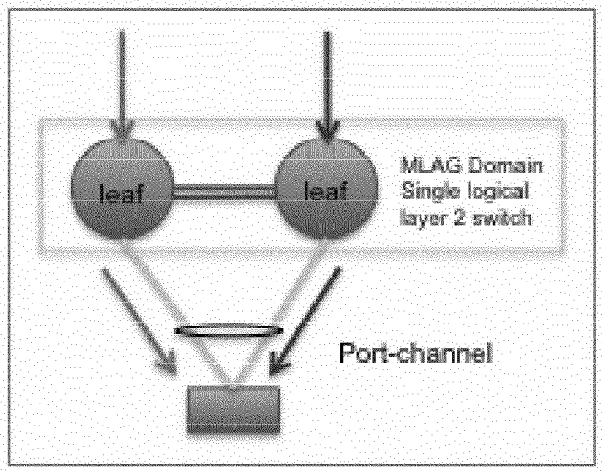

Figure 2. Traffic within the MLAG topology always flows via the optimal path

Spanning Tree

The MLAG technology removes any physical loops from the Layer 2 topology and the need therefore to deploy a spanning-tree protocol (STP, PVST and MSTP). Arista understands that in a Data Center deployment there will be a requirement to connect legacy devices. STP can therefore be enabled in conjunction with the MLAG, allowing standard STP interaction with legacy and third-party edge nodes while still maintaining an active-active Layer 2 topology.

Virtual-ARP

The MLAG technology, allows a pair of physical switches to be appear as a single logical Layer 2 switch. The two physical switches however maintain independent Layer 3 control plane. To provide an active-active first-hop gateway for the connected hosts, the technology also includes a virtual-ARP (VARP) functionality.

The VARP functionality allows the MLAG Domain to be configured with a Virtual ARP address, which would act as the first-hop default gateway for the attached hosts. Both physical switches within the MLAG Domain respond to ARP request to VARP address and route any subsequent traffic destined to the VARP MAC address. Thus allowing stateless Layer 3 forwarding of traffic through either physical switch to the upstream Spine layer. As the VARP functionality is stateless, there is no CPU overhead or scaling issues as the number of VARP instances increases.

The physical switches of the MLAG Domain maintain their independent Layer 3 control plane and therefore routing table. Thus in a Leaf-Spine architecture each physical switch would implement it's own independent IGP protocol and peer independently with the Spine. In a four- or eight-switch Spine model, this would mean each switch would independently peer with each Spine switch for load balancing and redundancy.

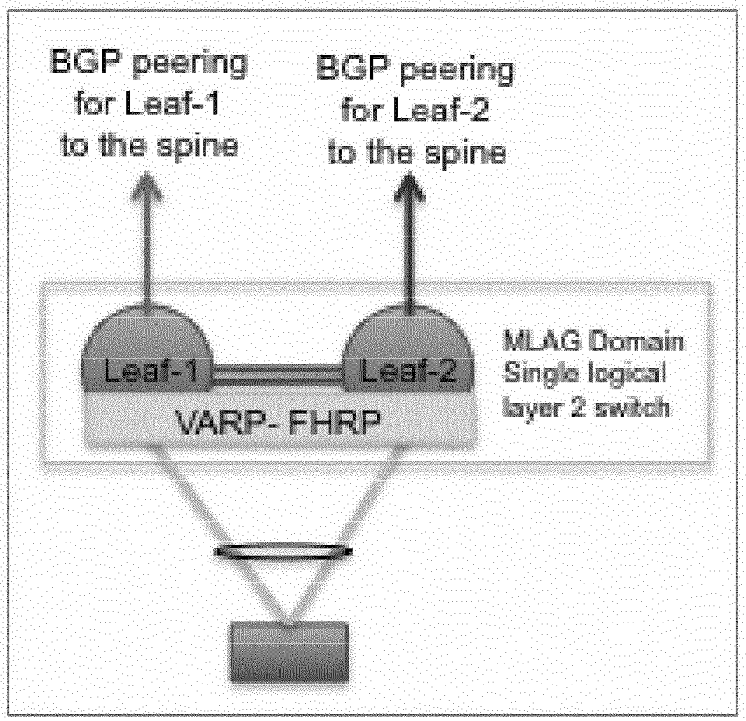

Figure 3. V-ARP Deployed as FHRP for Hosts with Leaf Running Individual BGP Peering to the Spine

ISSU and SSO

The physical Leaf switches within the MLAG Domain provide both Stateful Switch failover (SSO) and In-service-software Upgrades (ISSU) for the MLAG logical switch. In the event of a physical switch failure, the remaining active switch of the MLAG Domain, will continue to forwarded traffic at both Layer 2 and Layer 3 and if required seamless take ownership of the STP state machine without creating a re-convergence within the network. This SSO functionality is extended to provide ISSU for the MLAG Domain, with the ability to individually upgrade the physical switches within the domain sequentially with one switch always being active and forwarding traffic at any one time during the upgrade process.

1.4 Network Virtualization

The design has been constructed to provide a stable, high performance and scalable architecture. The adoption of virtualization and the resultant dynamic placement of customer services, means customers are no longer tied to a physical NTE switch, but rather spread across multiple NTEs within the IPAG architecture. This sprawl of resources and the need to provide L2 adjacency between them drives a requirement for network virtualization across the L3 fabric.

To achieve this network virtualization, Arista recommends industry-adopted standard, VXLAN technology. VXLAN is an IETF RFC co-authored by Arista's CTO Ken Duda, in conjunction with VMware, Citrix, Broadcom, and Red Hat. VXLAN allows the provisioning of L2 domains, called virtual network identifiers (VNIs), over a standard L3 architecture. This allows a customer's services to be L2 adjacent regardless of where they reside in the IPAG network.

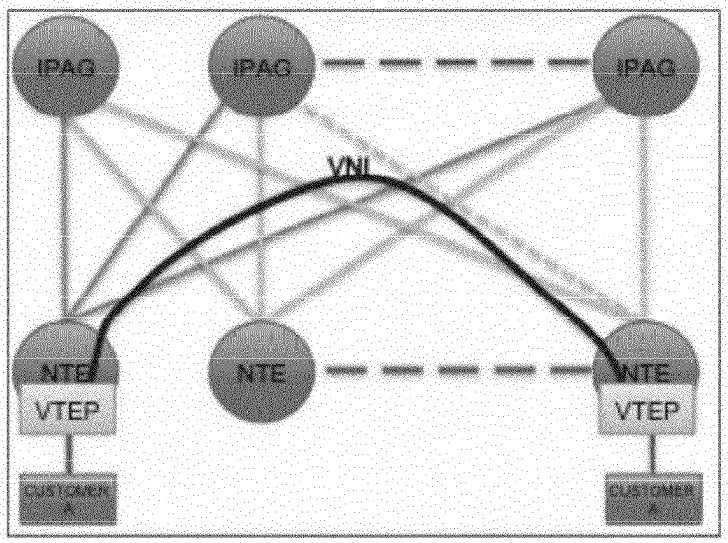

Figure 4. Layer 2 Adjacency between Customer A sites using VXLAN Encap/Decap across the Layer 3 IPAG

The VXLAN protocol operates by encapsulating the host traffic (MAC and IP header) within an outer IP UDP VXLAN header. The encapsulated packet is routed across the Leaf-Spine fabric to its destination based on the outer header irrespective of the encapsulated inner MAC and IP header. At the remote destination the packet is de-encapsulated and inner frame is forwarded to the end host.

1.4.1 Overview of the Arista VXLAN solution

The VXLAN encapsulation and de-capsulation process is achieved on a Virtual Tunnel End Point (VTEP) located on the Arista NTE switch at Switched Ethernet Services network.

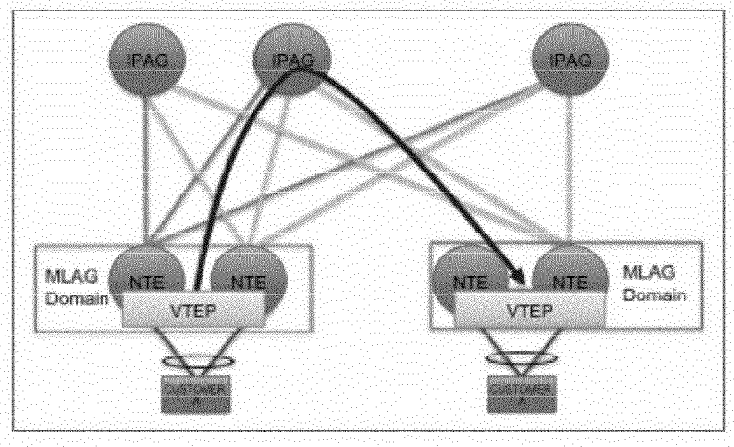

Figure 5. VXLAN in conjunction with MLAG for Layer 2 active-active connectivity into the VNI

The VXLAN functionality will provide L2 connectivity between customer sites across the L3 IPAG fabric. To provide resilient L2 connectivity for the clients, the VXLAN functionality will be combined with Arista's MLAG technology. This is achieved by creating a virtual VTEP across the MLAG pair, with the host as normal connecting to the MLAG Domain via a port-channel split across the two NTE switches of the MLAG Domain. Traffic from the client site can be load-balanced across both links of the port-channel, and VXLAN

encapsulated by either physical switch of the MLAG Domain. This provides active-active connection into the VXLAN domain, while removing any physical loop within the VNI. To achieve L3 connectivity between the VNIs, the Arista nodes will also provide VXLAN+Routing capability with VARP functionality for active-active L3 forwarding between the VNIs.

1.5 Arista's VXLAN Control plane

To flood unknown unicast traffic for MAC address learning and to forward broadcast traffic within a VNI, the VXLAN protocol uses IP multicast for the control plane. In this approach VTEP members of a VNI, join an associated IP multicast group and any broadcast or multicast traffic is forwarded to each of the VTEP members of the group. This approach places a requirement for IP multicast support on the L3 IPAG fabric, to remove this requirement Arista provides support for an additional control plane model which is termed head-end replication (HER).

The head-end replication (HER) approach allows the configuration of a remote VTEP flooded list for each VNI. Any broadcast, unknown unicast and multicast traffic (BUM) traffic for the VNI would be encapsulated by the VTEP and individually unicast to each of the remote VTEPs within the configured flood list. The functionality therefore removes any requirement for IP multicast support in the IPAG fabric for forwarding VXLAN BUM traffic. The programming of the flood-list can be a manually process or automated via the Arista eAPI by a AT&T's OAM platform or extension.

The MAC address-learning function in this model can be achieved using the normal flood and learn process or to reduce the level of flooding locally learned MAC addresses can be dynamically distributed to all remote VTEPs within the specific VNI. The dynamic distributed of the MAC addresses is achieved by enabling the CloudVision Extension (CVX) service on the EOS OS of an Arista switch. The EOS OS and the CVX service can also if required run on external virtual machine. The CVX process utilizes the distributed architecture of the EOS operating to mount the local VXLAN forward table of each of the switches, providing a global VXLAN forwarding table, which can then be mounted by the individual VTEPs.

1.6 Network Automation

Arista provides the most open and programmable network operating systems in the industry today, much of this has come from an early understanding of what is required to build a world class data center. A core function of this is the goal to reduce operational costs or OPEX. Some key components of our automation story focus on bare metal provisioning (ZTP). Arista also includes support for automation tools such as Chef, Puppet, CFengine, and Ansible.

A note about automation initiatives based on industry experience: investing the time and effort up front to properly build and test an automation solution will truly set AT&T on a path for success. Layering in automation to an already provisioned network can be a challenge. Some of the best automation successes occur when designed and built into the solution at initial deployment—by making automation a requirement from the start, implementation of the solution becomes ingrained in the network operators, managers and engineers. This type of bottoms-up approach has the most chance to break the boundaries of what has been traditionally possible with legacy network equipment.

1.7 Zero Touch Provisioning (ZTP)

Building off of standards based protocols (DHCP, TFTP, and HTTP) combined with advanced scripting support network administrators can now rapidly provision NTE switches in an automated fashion.

At boot time the switch will send DHCP requests out all interfaces, when received by the DHCP server, DHCP options in the form of http headers (X-DHCP-Option-…) sent by the Arista switch will identify the device and send other

information such as (System MAC, Hardware Version, Serial number, etc.) to the DHCP server. This information can then be used to automate the system based on unique system attributes. Once identified a boot script is passed back to the device with instructions on what to do, EOS version, interface or global configuration parameters can be set.

While a simple feature, the possibilities created by this are truly left up to the imagination. Several examples of what can be achieved are available. From full-on provisioning of the systems to self-configuring network devices, with EOS the possibilities are truly endless.

1.8 Custom Scripting (Beyond ZTP)

ZTP in its basic form is very useful for deploying and maintaining infrastructure but some want to take it further. The basic operation of ZTP can be further enhanced to provide advanced automation. For example, using LLDP (Link Level Discovery Protocol) to determine the location and role of the switch in the network. Once the role and location are determined from pre-defined criteria the device can auto insert into the network. This scripting can then download extensions, as well as perform a variety of other useful tasks.

Arista is committed to helping AT&T take full advantage of all the automation capabilities of our platforms. Our customer engineering teams, working in conjunction with your automation teams and partners, will design an automation solution that meets AT&T requirements. This EOS+ service will be included with our proposal to help install a custom solution to fit AT&T's requirements in a turnkey manner.

At Arista we take great pride in contributing to the open nature of our EOS network operating system. We have a powerful open development community (eos.arista.com) where users can share find tech tips, automation scripts, and share experiences on Arista devices. The community approach allows us to share information and create an ecosystem of user-driven open tools. Many of these tools and tech tips, have a common theme of automation, proactive notification, and true differentiation.

1.9 Arista API (eAPI)

Arista has developed an open northbound API for issuing commands to the switch or group of switches, the API known as EAPI is based on JSON RPC over http/https. EOS API (eAPI) is a comprehensive tool for network managers. eAPI is an external API through which applications or controllers can gain visibility into system state. It allows customers to integrate applications and other network components with Arista's switching platforms. eAPI provides a mechanism to run legacy applications seamlessly against newer versions of EOS.

eAPI is protocol and data model agnostic – allowing the customer to choose the most familiar programming language for third party integration

Some of the advantages of our eAPI's JSON structure:

- Less verbose
- More careful in it's use of data types
- Deeply interoperable with any programming language
- Ability to represent most general data structures: records, lists and trees
- Significantly simpler syntax – more efficient parsing
- Can represent any non-recurrent data structure as is – more flexible

Only a scalable API such as eAPI can support true network agility. eAPI is complete and portable. It extends across various layers of the network from Leaf to Spine. eAPI provides a shared mechanism to support CLI, SNMP, and

programming APIs under a common scheme. eAPI provides a seamless integration point for controller or applications to read and modify the state of the switch within the network—thus fully bringing the software-defined networking concept to life.

1.10 Automation Tools (Chef, Puppet, CFengine, and Ansible)

Arista EOS supports native integration with a variety of network orchestration and configuration management tools. The same tools used by system administrators to automate provisioning across large numbers of servers can also be used to orchestrate Arista's switch portfolio. Arista's open and extensible operating system provides unparalleled support for off the shelf Linux tools. Reference configuration, as well as code examples are published on our EOS Central website as well as GitHub for general deployment and use.

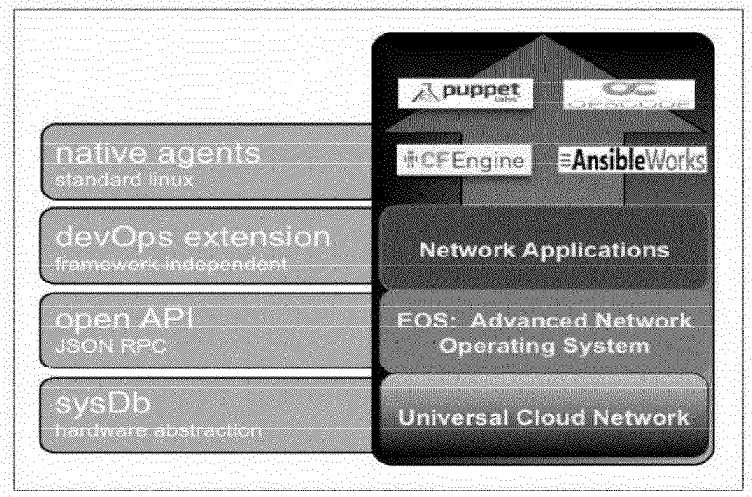

Figure 6. Arista Open APIs for Integration with Third-party Tools and AT&T In-house Deployment Tools

2 Arista Networks

2.1 About Arista Networks

Arista Networks was founded to deliver software driven cloud networking solutions for large data center and high-performance computing environments. With more than two million cloud networking ports being deployed worldwide, Arista delivers a portfolio of 1/10/40 and 100GbE products that redefine network architectures, bring extensibility to networking, and dramatically change the price/performance of data center networks.

At the core of Arista's platform is the Extensible Operating System (EOS™), a groundbreaking network operating system with single-image consistency across hardware platforms and modern core architecture enabling in-service upgrades and application extensibility.

Arista is recognized by Gartner as a "visionary" in its 2014 "Magic Quadrant for Data Center Networking" based on its ability to execute and completeness of vision. The Arista team is comprised of experienced management and engineering talent from leading networking companies. Arista designs revolutionary products in California and delivers them worldwide through distribution partners, systems integrators, and resellers with a strong dedication to partner and customer success.

2.2 Strong Management Team

Arista was founded in 2004, launched in 2008 and is led by Jayshree Ullal. The seasoned management team has a rich and extensive history in networking and innovation. The Arista team is globally recognized as respected leaders and visionaries in networking technology.

2.2.1 Jayshree Ullal: President and Chief Executive Officer

Jayshree Ullal is a networking executive veteran with 25+ years of experience and was named one of the "50 Most Powerful People" in 2005 Network World and one of the "Top Ten Executives" at VMworld 2011. As President and CEO of Arista Networks, she is responsible for building the company's business in cloud networking.

Formerly, Jayshree was Senior Vice President at Cisco and responsible for $10B in annual revenue from Data Center, Switching and Services, including Cisco's flagship Nexus 7000 and Catalyst 4500 and 6500 product lines. During her tenure at Cisco, Jayshree forged key alliances with EMC, VMware and Microsoft in virtualization and application acceleration. Prior to joining Cisco, Ullal was the Vice President of Marketing at Crescendo Communications, which was Cisco's first acquisition in 1993.

Jayshree holds a B.S. in Engineering (Electrical) from San Francisco State University and an M.S. degree in engineering management from Santa Clara University.

2.2.2 Andy Bechtolsheim: Founder, Chief Development Officer and Chairman

As Chief Development Officer, Andy Bechtolsheim is responsible for the overall product development and technical direction of Arista Networks.

Previously Andy was a Founder and Chief System Architect at Sun Microsystems, where most recently he was responsible for industry standard server architecture. Andy was also a Founder and President of Granite

Systems, a Gigabit Ethernet startup acquired by Cisco Systems in 1996. From 1996 until 2003 Andy served as VP/GM of the Gigabit Systems Business Unit at Cisco that developed the very successful Catalyst 4500 family of switches. Andy was also a Founder and President of Kealia, a next generation server company acquired by Sun in 2004.

Andy received an M.S. in Computer Engineering from Carnegie Mellon University in 1976 and was a Ph.D. Student at Stanford University from 1977 until 1982.

2.2.3 Kenneth Duda: Founder, CTO and Senior Vice President, Software Engineering

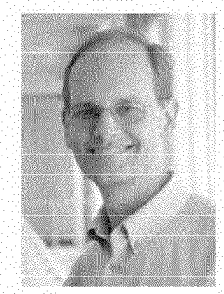

Kenneth Duda is a pioneer in high-performance networking software and lead architect of Arista Networks EOS, a stateful modular operating system for all Arista Networks products. He is also the co-author of network virtualization specifications including VXLAN with VMware and NVGRE with Microsoft. From 2005 to 2008, Ken was also the Acting President of Arista Networks.

Prior to joining Arista Networks, Ken was the CTO at There.com, where he played a lead role in designing a real-time 3-D distributed system that scaled to thousands of simultaneous users. Ken was also the first employee of Granite Systems and led the software development effort for the Catalyst 4000 product line after the acquisition by Cisco.

Ken has 3 simultaneous engineering degrees from MIT and holds a Ph.D. in Computer Science from Stanford University.

2.2.4 Anshul Sadana: Senior Vice President, Customer Engineering

Anshul Sadana has over a decade of experience in engineering management, design & software development in the networking industry. Anshul joined Arista in 2007 and is responsible for product definition and development priorities of Arista's next generation products. He also leads all systems engineering and customer advocacy functions.

Previously at Cisco, Anshul led the development of the Catalyst 4500 and 4900 product lines and also managed strategic customer relations.

Anshul has an M.S. in Computer Science from the University of Illinois and an M.B.A. from the Wharton School of Business.

2.3 Business and Technical Points of Contact

2.3.1 AT&T Account Team

For future correspondence related to this proposal, AT&T may contact the following individuals or the broader AT&T team at ATT@Arista.com.

Business Contact

Wayne Stickle
Director of Sales, AT&T
Arista Networks, Inc.
5453 Great America Parkway
Santa Clara, CA 95054

Office: (512) 263-0571
Mobile: (512) 656-6373

Email: Wayne@Arista.com

Technical Contact

Steve McCarthy
Director of Technology, AT&T
Arista Networks, Inc.
5453 Great America Parkway
Santa Clara, CA 95054

Office/mobile: (703) 851-2855

Email: McCarthy@Arista.com

2.3.2 Global Engagement Model

Arista has a global and local engagement model with AT&T. Details of global team and roles are as follows.

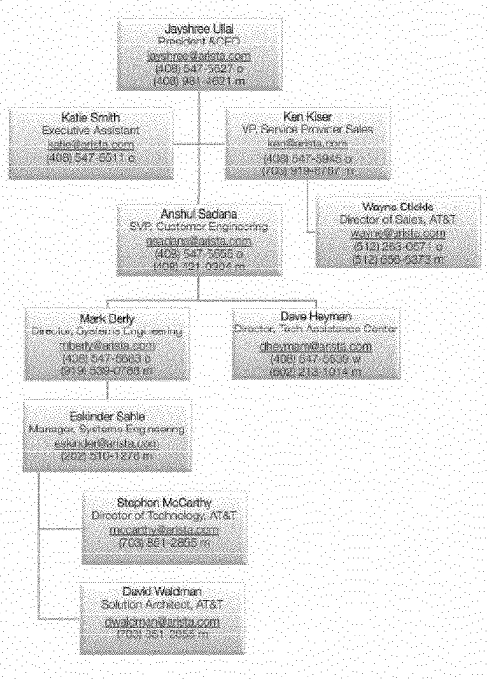

Figure 7. AT&T Global Account Team Structure

2.4 Arista Product Offerings

2.4.1 The World's Most Advanced Network Operating System

Arista EOS is the core of Arista cloud networking solutions for next-generation data centers and cloud networks. Cloud architectures built with Arista EOS scale to tens of thousands of compute and storage nodes with management and provisioning capabilities that work at scale. Through its programmability, EOS enables a set of software applications that deliver workflow automation, high availability, unprecedented network

visibility and analytics and rapid integration with a wide range of third-party applications for virtualization, management, automation, and orchestration services.

Arista Extensible Operating System (EOS) is a fully programmable and highly modular, Linux-based network operation system, using familiar industry standard CLI and runs a single binary software image across the Arista switching family. Architected for resiliency and programmability, EOS has a unique multi-process state sharing architecture that separates state information and packet forwarding from protocol processing and application logic.

Video 1. Arista Powers the Cloud with EOS (3:22)

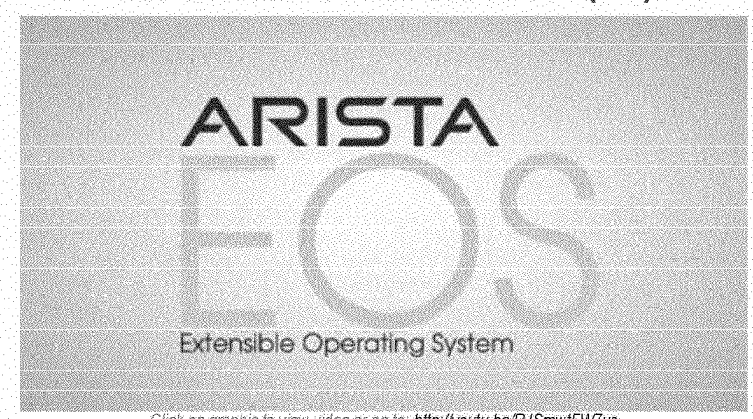

Click on graphic to view video or go to: http://youtu.be/RJSmwtFW7vs

Cloud Scale Architecture

- Scale to your needs, from 100 to 100,000+ compute and storage nodes
- Rich management and provisioning capabilities that work at scale
- Open, standards-based approach with MLAG at Layer 2, ECMP at Layer 3 with effective use of all available bandwidth in non-blocking modes while providing failover and resiliency
- Network virtualization using tunneling technologies such as VXLAN and NVGRE for seamless workload mobility
- Innovative Spline™ architecture for high density hosts in a single-tier or two-tier network

Open and Programmable

- Open integration with all application and infrastructure elements via eAPI and Advanced Event Manager (AEM)
- Programmable at all layers: Linux kernel, hardware forwarding tables, Virtual Machine orchestration, switch configuration, provisioning automation and advanced monitoring
- EOS Application Extensibility for the ability to run cloud infrastructure automation applications (such as Chef, Puppet or Ansible) and network analytics applications (such as Splunk)
- Easily adaptable to in-house network management systems
- Key enabler of software applications for automation and visibility, such as Zero Touch Provisioning, VM Tracer, and Latency Analyzer (LANZ). Learn more at Arista.com

High Availability

- Reduce maintenance windows with Arista Smart System Upgrade (SSU) reduces maintenance windows through intelligent insertion and removal of network elements
- In-service software upgrade (ISSU) for individual processes within EOS

- Self-healing resiliency for minimum downtime with fault containment to a single module and process restart without the need to rebuild state information
- Custom monitoring, failover and load balancing with third-party integration for custom monitoring, failover and load balancing

Visibility

- Unprecedented visibility into application performance and network-wide monitoring capabilities for both industry standards and customer specific dev/ops solutions
- Simplified Tap Aggregation with the Arista Data Analyzer (DANZ) feature set
- Rapid identification and troubleshooting of application and network performance problems through tracers such as VM Tracer, Latency Analyzer (LANZ), MapReduce Tracer, sFlow, and Path Tracer

Automation

- Simplified provisioning for new and replacement switches with Zero Touch Provisioning (ZTP) and Zero Touch Replacement (ZTR)
- Advanced Event Management for automated responses to network and application events
- Automate complex IT workflows and simplify network operations to individual requirements through rich programmatic capabilities
- Automation Integration with partners enhance native capabilities with tools such as Puppet, Chef and Ansible and extends automation up the stack to include other network systems and applications including firewalls, load balancers and compute infrastructures with partners such as F5, VMware/Nicira, and Nuage.

2.4.2 Arista Networks Cloud Networking Portfolio

Arista Networks is the leader in building scalable high-performance and ultra-low-latency networks for today's data center and cloud computing environments.

Purpose-built hardware, and Arista EOS, the world's most advanced network operating system, provide single-binary system images across all platforms, maximum system uptime, stateful fault repair, Zero Touch Provisioning, Latency Analysis and a fully accessible Linux Shell. Arista Ethernet switches are the perfect network solution for your most demanding workloads. With native support for VMware Virtualization and hundreds of Linux applications integrated into hardware platforms designed to meet the stringent power and cooling requirements of today's most demanding data centers, Arista delivers the most efficient and best performing 10Gb Ethernet platforms.

For more details, please click here: Arista Products Quick Reference Guide.

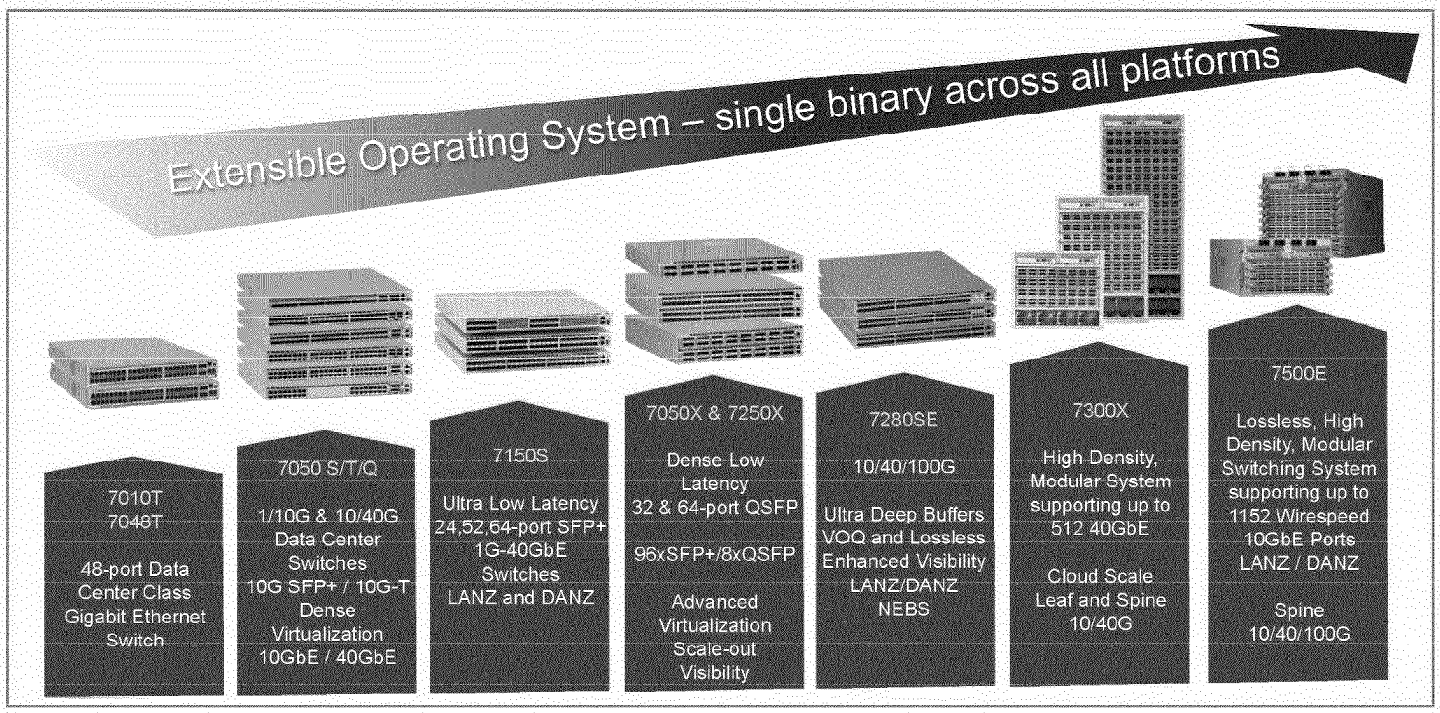

Figure 8. Arista Data Center Switches

2.5 Training

Often customers are hesitant to deploy an infrastructure based on a new vendor's solution due to the cost of training an entire 'army' of engineers to design, implement and operate a new vendor solution. Simply put Arista's EOS uses standard industry CLI for it's switching platforms. As such, very often an existing customer can 'copy & paste' the configurations from their existing Cisco infrastructure onto an Arista switch and 90%+ of commands will be accepted. Arista has gone to extraordinary lengths to ensure that all configurations and troubleshooting commands follow industry CLI standards. All this reduces to a minimum the need to re-train existing engineers for the migration to an Arista solution.

2.6 Support and Operations

Arista TAC engineers are on call 24x7x365 by phone or email to support AT&T regardless of the time zone in which your needs arise. Online access is available to allow you to create and review cases, download software, upload capture files or to download support information such as release notes. Worldwide RMA depots ensure multiple hardware replacement options, from "return-to-factory" with 10-business day replacement to "advanced" with hardware replacement options from same-day 2- or 4-hour shipment.

2.6.1 Contacting Support

There are multiple ways to get assistance from the Arista Networks Support organization:

- Phone:
  - US TollFree: +1-866-476-0000
  - US Toll: +1-408-547-5502

- UK TollFree: +44-808-234-0722
- UK Toll: +44-20-7023-9352

- Email: support@aristanetworks.com
- Online Tools: http://www.arista.com/en/support/customer-portal/tac-portal then Support --> Case Management

2.6.2 Arista Websites

Support Page

Arista's Customer Support webpage (http://www.aristanetworks.com/en/support) links to Software Downloads, EOS Documentation, Knowledge Base, Security Advisories and Field Notices.

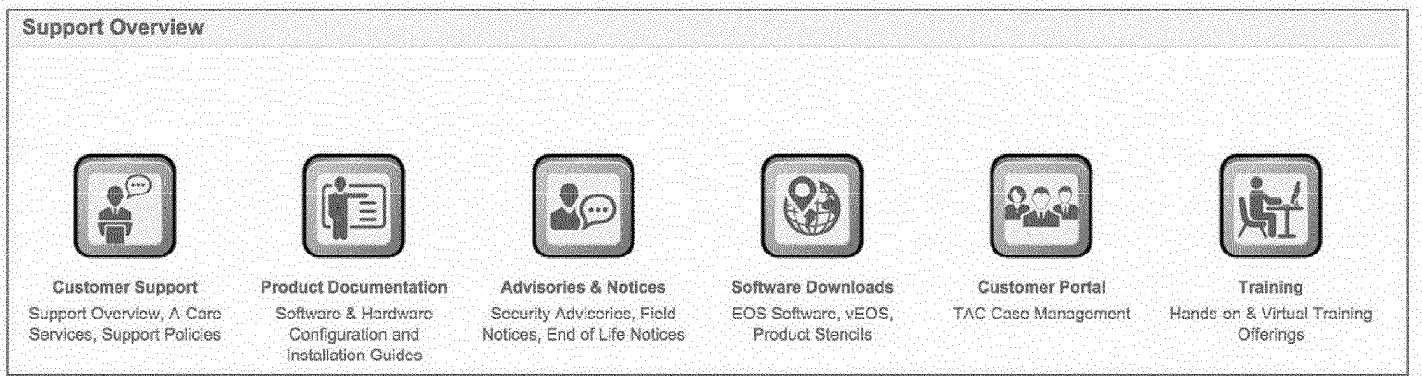

Figure 9. Arista Support Web Portal

EOS Central

Arista has a community support site (https://eos.aristanetworks.com/home.php). Here you can post general questions to the Forum, access to Tech Tips and scripts written by customers and by Arista employees, plus much more. Customer participation in this community is encouraged.

GitHub

Software repository on Github, for downloading custom scripts and EOS extensions.

2.6.3 Service Requests

A Service Request (SR) is opened when contacting Arista's global TAC organization. You will be assigned a case number in the format of *SR <5 digit number>*. This number should be referenced when working with TAC on a particular case.

A *show tech* is typically required for each SR. Beginning with 4.8 code, an hourly snapshot of the *show tech* is saved in flash:schedule/tech-support/. Providing a *show tech* before and after the problem should help the Customer Support Engineer (CSE) determine root cause.

Severity Levels

Arista provides the following guidance as to the severity of TAC cases, though our priority is to minimize the impact to your business regardless of the issue.

- **P1 – Mission critical network is down and/or severe degradation of the network.** Highest priority for a case. Customer Support Engineering, Sales team, Engineering team, and Executives expected to be involved.
- **P2 – Network deterioration is impacting the business.** High priority case. Indicates a need for urgency in resolution of the case.
- **P3 – Network performance is not ideal, but normal operation for business continues**. Default priority of all cases.
- **P4 – Zero network impact. Request is regarding configuration, installation, and general product questions.** Low priority case.

Status Levels

Below is a list of the various statuses that are applied to a SR.

- New - Case has been opened, but not assigned to a CSE.
- Assigned - Case assigned to a CSE and being worked on.
- Waiting-on - Assigned, but CSE is waiting on an action by someone else.
- Resolved - Case has been closed.
- Re-opened – It was previously determined that the case was successfully resolved, but the issue has again been experienced in the network.

2.6.4 Arista A-Care Offerings

Arista A-Care service offerings are designed to provide AT&T with a world-class support when needed. Arista TAC engineers leverage years of industry experience and can complement the AT&T engineering team with expertise on all Arista products. The A-Care service offerings are available 24x7x365 with advance replacement options to minimize any network downtime. All A-Care Service options include full access to bug fixes and software download

Table 1. Arista A-Care Support Services

	A-Care Next Business	A-Care 4 Hour	A-Care 2 Hour
Unlimited, 24x7 TAC access			
Software Download			
Online Case Management			
Arista Networks community forums			
Advance Replacement of hardware			
RMA Service level	Next Business Day*	4-Hour*	2-Hour*

** Delivery service levels are available to most metropolitan areas worldwide. Where not available, replacement product will be shipped the same day from the closest RMA depot, but actual delivery times may vary.*
Standard Limited Hardware Warranty: 1 year return to factory
Standard Limited Software Warranty: 90 days

2.6.5 Return Merchandise Authorization (RMA)

Contacting support for an RMA should be requested through a Service Request. If it is determined that there is faulty hardware the CSE will begin the RMA process. The speed in which the replacement hardware is received will be determined by the support contract in effect. The CSE can access the support contract information and inform you as to what level of service applies to the switch in question.

Worldwide RMA depots ensure multiple hardware replacement options, from "return-to-factory" with 10-business day replacement to "advanced" with hardware replacement options from same-day 2- or 4-hour shipment. Arista Networks is able to quickly provision new RMA depots, should they be required, to facilitate AT&T global requirements

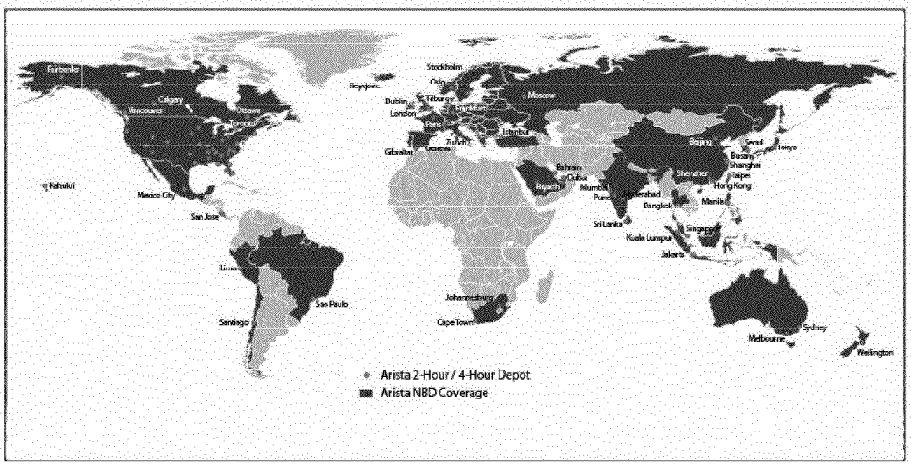

Figure 10. Worldwide RMA Depot Locations

2.6.6 Escalation Procedure

Standard escalation can be handled through the use of the SR Priority. As previously mentioned, higher priority cases have increase visibility within Arista. Typically the use of the SR Priority provides the customer with the necessary escalation required.

In extreme situations, when further escalation becomes necessary, the Arista account team can be contacted directly and broadly using the following email distribution address: ATT@Arista.com.

2.6.7 EOS Life Cycle Policy

Arista Networks has established an EOS Software Release Policy and life cycle guideline to help customers and partners facilitate EOS migration and plan multi-year infrastructure deployment. To assist AT&T in selecting the right EOS software release for their environments, Arista has adopted a naming convention of EOS releases. The naming convention identifies if a particular release is integrating new feature functionality, or has reached software maintenance mode.

Starting with EOS 4.13, each EOS software release will be identified with either the letter "***F***" indicating that it contains new *F*unctionality release or "***M***", denoting that the release is in the *M*aintenance phase. An **M** release only receives incremental fixes and no new functionality is added.

Arista EOS software release numbering will follow the following example:

4.X.0F : First release
4.X.1F : Next Feature Release
4.X.2F : Addition of new features as required

4.X.YM : First release containing bug fixes only followed by ongoing maintenance
4.X.ZM : Last maintenance release

Arista will support each major EOS software release for up to 30 months from the date of FCS (First Customer Ship).

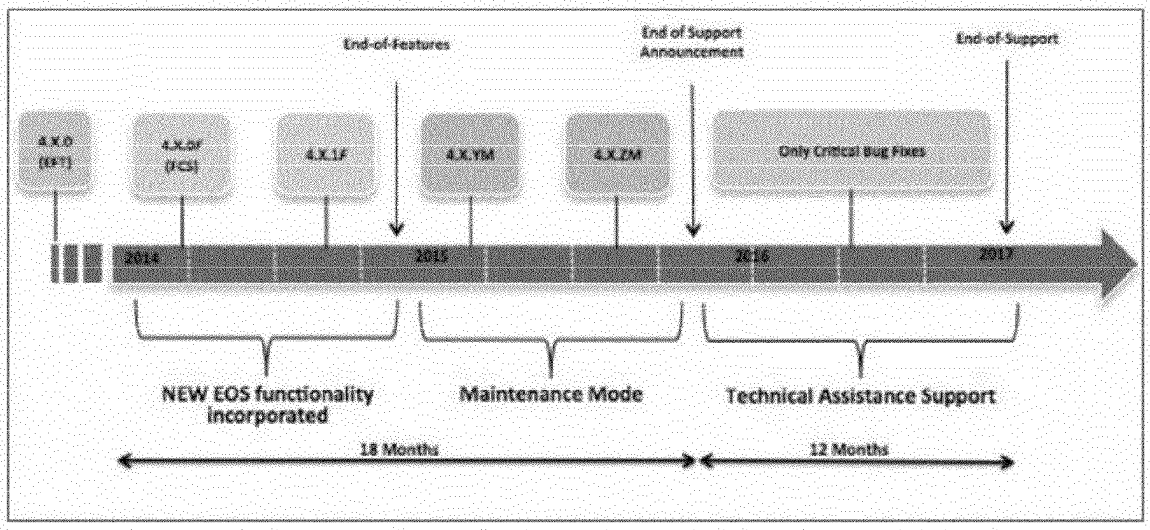

Figure 11. EOS Software Release Cycle

For more information about the Arista EOS software end of life policy

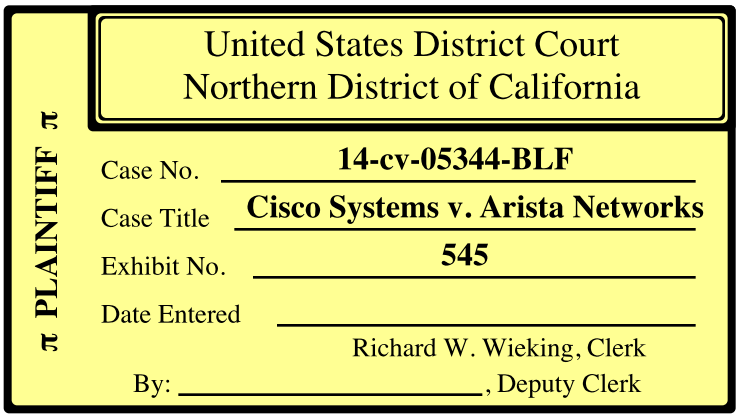
United States District Court
Northern District of California
PLAINTIFF
Case No. 14-cv-05344-BLF
Case Title Cisco Systems v. Arista Networks
Exhibit No. 545
Date Entered
Richard W. Wieking, Clerk
By: , Deputy Clerk